

1 METHODS AND APPARATUS FOR SUPPORTING
2 MULTIPLE UTOPIA MASTERS ON THE SAME UTOPIA BUS

3
4 BACKGROUND OF THE INVENTION

5
6 1. Field of the Invention

7 The present invention relates to telecommunications. More
8 particularly, the present invention relates to the passing of high
9 speed Asynchronous Transfer Mode (ATM) data over a standardized
10 Universal Test and Operations Physical Interface for ATM (UTOPIA)
11 bus.

12
13 2. State of the Art

14 Perhaps the most awaited, and now fastest growing technology
15 in the field of telecommunications in the last decade is known as
16 ATM (Asynchronous Transfer Mode) technology. ATM is providing a
17 mechanism for removing performance limitations of local area
18 networks (LANs) and wide area networks (WANs) and providing data
19 transfers at a speed on the order of gigabits/second. Within the
20 ATM technology, a commonly used interface specification between
21 chips on a board for passing ATM cells is the UTOPIA (Universal
22 Test & Operations PHY Interface for ATM) interface. The UTOPIA
23 interface is specified in ATM Forum standard specifications: af-
24 phy-0017.000 (UTOPIA Level 1, Version 2.01 March 21, 1994);

1 af_phy_0039.000 (UTOPIA Level 2, Version 1, June 1995); and af-
2 phy-00136.000 (UTOPIA 3 Physical Layer Interface November 1999)
3 which are hereby incorporated by reference herein in their
4 entireties.

5
6 As mentioned above, the UTOPIA interface is defined by the
7 ATM Forum to allow a common interface between the Physical Layer
8 (PHY) and ATM layer in an ATM system. Currently, four levels of
9 UTOPIA interfaces are defined by the ATM Forum to support a wide
10 range of speed and media types from low speed xDSL to high speed
11 OC-192. A typical application of the UTOPIA interface is
12 supporting the connection between an ATM network processor and
13 various PHY devices such as a DSL chip set and/or a SONET framer.
14 UTOPIA may also used as the interface between a switch fabric and
15 an ATM network processor.

16
17 UTOPIA supports three operation modes: single PHY operation
18 mode, Multiple PHY (MPHY) with Direct Status Indication operation
19 mode and MPHY with Multiplexed Status Polling operation mode. In
20 the single PHY mode, the UTOPIA interface includes a data bus and
21 a control bus. The operation of UTOPIA in the single PHY mode is
22 relatively simple and straightforward. In MPHY operation mode,
23 the UTOPIA interface includes a data bus, a control bus and an
24 address bus.

1 The MPHY UTOPIA transmit interface includes the following
2 signals: transmit data (TxData), transmit address (TxAddr), and
3 the transmit control signals including transmit cell available
4 (TxClav), transmit enable (TxEnb*) and transmit start of cell
5 (TxSOC). The receive interface includes the following signals:
6 receive data (RxData), receive address (RxAddr), and the receive
7 control signals including receive cell available (RxClav), receive
8 enable (RxEnb*) and receive start of cell (RxSOC). A MPHY device
9 may consist of multiple logical PHY ports. Each PHY port has a
10 one-to-one correspondence with a PHY Port address that is related
11 to a UTOPIA address and Clav (Cell buffer available) signal.

12
13 Prior art Figure 1 illustrates an example of a UTOPIA Level 2
14 interface supporting MPHY with Multiplexed Status Polling
15 operation. Multiplexed Status Polling is mainly used in actual
16 applications. As shown in Figure 1, a transmit clock signal
17 (TxClk) is used to clock control signals and data signals in the
18 transmit direction (from the ATM device to the PHY devices). The
19 TxData[15:0] signal is a 16-bit UTOPIA transmit data bus. The
20 assertion of TxEnb* is coincident with the start of the cell
21 transfer. TxSOC is used to indicate the start of cell position.
22 TxClav is used to indicate that the PHY layer device is ready to
23 receive a cell from the ATM layer device. TxAddr[4:0] is the

1 UTOPIA address and is used to poll and select the appropriate MPHY
2 device.

3
4 At the UTOPIA transmit interface, the ATM layer device polls
5 the TxClav status of a PHY layer device by placing a specified
6 address on the TxAddr bus for one clock cycle. The PHY layer
7 device which is associated with the address on the TxAddr bus
8 drives TxClav high (or low) during the next clock cycle during
9 which the ATM device places a null address (1F) on the TxAddr bus.
10 The ATM layer device checks TxClav at a certain time after it
11 issues TxAddr. Based on polled TxClav information, the ATM layer
12 device can select a PHY device and transfer data to this PHY
13 device by driving TxEnb*, TxAddr, and TxSOC signals.

14
15 Similarly, RxClk is the receive clock signal that is used to
16 clock control signals and data in the receive direction (from the
17 PHY device to the ATM device). RxData[15:0] is a 16-bit UTOPIA
18 Receive bus. The assertion of RxEnb* is coincident with the start
19 of the cell transfer. RxSOC is used to indicate the start of cell
20 position. RxClav is used to indicate that the PHY layer device is
21 ready to Receive a cell from the ATM layer device. RxAddr[4:0] is
22 the UTOPIA address of the PHY device and is used by the ATM device
23 to poll and select the appropriate PHY device in the receive
24 direction.

1 At the UTOPIA receive interface, the ATM layer device polls
2 the RxClav status of a PHY layer device by placing a specified
3 address on RxAddr bus for one clock cycle. The PHY layer device
4 which is associated with the address on the RxAddr bus drives
5 RxClav high (or low) during the next clock cycle during which the
6 ATM device places a null address (1F) on the RxAddr bus. The ATM
7 layer device checks RxClav at a certain time after it issues
8 RxAddr. Based on polled RxClav information, the ATM layer device
9 can select a PHY device and receive data from this PHY device by
10 driving RxEnb* and RxAddr signals.

11
12 Certain timing requirements must be met for the Multiplexed
13 Status Polling operation of the UTOPIA interface so that the ATM
14 layer device can correctly detect Clav (Cell buffer available)
15 information. Figure 2 depicts the timing requirement for UTOPIA
16 level 2 address polling. An ATM device starts driving UTOPIA
17 address N at time t1. The PHY device having address N responds to
18 address polling by driving the Clav signal high (or low) at time
19 t2. The Clav becomes valid for the ATM device at time t3 (after
20 transmission delay) as illustrated by the last line of the timing
21 diagram. The ATM device checks the Clav signal at time t4 in
22 order to accommodate a certain amount of transmission delay. To
23 guarantee correct operation, the Clav signal must be valid before

1 it is checked by ATM device, i.e. at some $t_3 < t_4$. This is a
2 necessary timing requirement for the UTOPIA interface.

3
4 From the foregoing, it will be appreciated that UTOPIA allows
5 many PHY devices to communicate with one ATM layer device. Access
6 to the UTOPIA bus is controlled by the ATM layer device which is
7 considered the bus "master". The PHY devices are thus considered
8 bus "slaves". It would be desirable, however, to allow more than
9 one ATM layer device to be coupled to the same UTOPIA bus so as to
10 provide either redundancy or load sharing.

11
12 SUMMARY OF THE INVENTION

13
14 It is therefore an object of the invention to provide methods
15 and apparatus for supporting multiple UTOPIA masters on the same
16 UTOPIA bus.

17
18 It is also an object of the invention to provide methods by
19 which multiple ATM layer devices can share a single UTOPIA bus.

20 It is another object of the invention to provide methods and
21 apparatus for supporting multiple UTOPIA masters on the same
22 UTOPIA bus which do not require any modification of the standard
23 UTOPIA bus.

1 It is still another object of the invention to provide ATM
2 layer devices which communicate outside the UTOPIA bus in order to
3 share control of the UTOPIA bus.

4
5 In accord with these objects which will be discussed in
6 detail below, the methods of the present invention include
7 coupling two UTOPIA bus masters via three signal lines (Ready,
8 Request, and Grant), designating one of the masters a primary
9 master and the other a secondary master, and coupling both bus
10 masters to the same UTOPIA bus. When receiving cells from the
11 UTOPIA bus, the secondary master asserts the Ready line when its
12 buffer has room to accept a cell. The primary master will not
13 assert the UTOPIA RxEnb unless the Ready line is asserted. When
14 RxEnb is asserted, both masters accept cells and screen them (with
15 a lookup table) to determine which cells are addressed to them.
16 When transmitting cells, the primary master normally controls
17 polling and PHY selection. The secondary master asserts the
18 Request line during polling to indicate that it has a cell to send
19 to a PHY that has responded positively to the polling. In
20 response to the Request signal, the primary master asserts the
21 Grant line and control of the bus is given to the secondary master
22 for the next cell cycle. According to one embodiment, if the
23 secondary master fails to assert the Ready line for more than one
24 cell time, it is assumed to be malfunctioning and is ignored until

1 it is reset. The presently preferred embodiment utilizes an
2 arbitration scheme to assure fairness in allocating control of the
3 bus. The presently preferred arbitration scheme requires that
4 both bus masters maintain a "scoreboard" of all PHYs indicating
5 (with one bit for each PHY) the results of the most recent poll
6 (i.e. which PHYs asserted the CLAV line). Each time a PHY is
7 selected by either master, the corresponding CLAV bit in the
8 scoreboard is reset. Neither master may select a PHY unless the
9 corresponding CLAV bit is set in the scoreboard. According to the
10 presently preferred embodiment, when control of the bus is
11 relinquished by either master, a dead cycle follows during which
12 neither master controls the bus and a weak pullup maintains all
13 signals in the high state.

14

15 The apparatus of the invention includes a modified ATM layer
16 device having a UTOPIA interface and means for coupling it via
17 Ready, Request, and Grant lines to another such device. The
18 modified device also includes logic for carrying out the methods
19 of the invention. The presently preferred apparatus is based on
20 the Transwitch® Cubit® Multi-PHY CellBus® Access Device. A single
21 device can be either primary or secondary master. The designation
22 is made with a lead strap.

1 Additional objects and advantages of the invention will
2 become apparent to those skilled in the art upon reference to the
3 detailed description taken in conjunction with the provided
4 figures.

5
6 BRIEF DESCRIPTION OF THE DRAWINGS

7
8 Figure 1 is an illustration of a prior art UTOPIA ATM
9 interface;

10
11 Figure 2 is a simplified prior art timing diagram
12 illustrating the polling of PHY layer devices by an ATM layer
13 device;

14
15 Figure 3 is a simplified block diagram illustrating an
16 embodiment of the invention in conjunction with a UTOPIA bus and a
17 plurality of PHY devices;

18
19 Figure 4a is a timing diagram illustrating the UTOPIA
20 selection process when the primary master is in control of the
21 bus;

1 Figure 4b is a timing diagram illustrating transfer of
2 control to the secondary master and the UTOPIA selection process
3 by the secondary master;

4
5 Figure 5 is a simplified block diagram illustrating a
6 presently preferred embodiment of the invention; and

7
8 Figure 6 is a simplified block diagram of an exemplary
9 application of the invention.

10
11 BRIEF DESCRIPTION OF THE APPENDIX

12
13 The Compact Disk appendix contains a 120 page data sheet
14 describing the presently preferred embodiment of the invention.
15 The Compact Disk is formatted for a Macintosh® computer and
16 contains a single file in Adobe® Acrobat® format.

17
18 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

19
20 Referring now to Figure 3, an arrangement 10 according to the
21 invention includes a first bus master 12 and a second bus master
22 14, both coupled to the same UTOPIA bus 16 which is coupled to
23 multiple PHY devices 18a, 18b, ..., 18n. The first bus master 12
24 is coupled to an ATM layer device or network 20 and the second bus

1 master 14 is coupled to an ATM layer device or network 22.
2 According to the invention, the two bus masters 12, 14 are coupled
3 to each other by three signal lines, Ready 24, Request 26, and
4 Grant 28. In addition, one of the bus masters is designated the
5 "primary" bus master and the other is designated the "secondary"
6 bus master. As shown in Figure 3, bus master 12 is the primary
7 bus master and bus master 14 is the secondary bus master. The
8 arrows on the signal lines 24, 26, and 28 reflect the relationship
9 between primary and secondary bus masters. The secondary bus
10 master 14 asserts the ready 24 and request 26 signals and the
11 primary bus master 12 asserts the grant 28 signal.

12
13 According to the methods of the invention, when receiving
14 cells from the UTOPIA bus 16, the secondary master 14 asserts the
15 Ready line 24 when its buffer has room to accept a cell. The
16 primary master 12 will not assert the UTOPIA RxEnb unless the
17 Ready line 24 is asserted. When RxEnb is asserted, both masters
18 12 and 14 accept cells and screen them (with a lookup table
19 described in more detail below with reference to Figure 5) to
20 determine which cells should be forwarded and which should be
21 ignored. According to one embodiment, if the secondary master 14
22 fails to assert the Ready line 24 for more than one cell time
23 (approximately one microsecond, it is assumed to be malfunctioning
24 and is ignored until it is reset. This is a fail-safe mechanism

1 so that a malfunction in the secondary master does not halt the
2 operation of the primary master.

3
4 When transmitting cells to the UTOPIA bus 16, the primary
5 master 12 normally controls polling and PHY selection. This is
6 illustrated in the timing diagram of Figure 4a. For example, the
7 primary master polls an address at time t_0 . The null address (1F
8 is placed on the address bus during the next clock cycle during
9 which the polled PHY device replies. A PHY device is selected in
10 the following clock cycle starting at time t_1 .

11
12 When the secondary master has a cell to transmit, it asserts
13 the Request line 26 during polling to indicate that it has a cell
14 to send to a PHY that has responded positively to the polling. In
15 response to the Request signal, the primary master asserts the
16 Grant line 28 and control of the bus 16 is given to the secondary
17 master for the next cell cycle. This is illustrated in Figure 4b
18 where the grant is issued at time t_2 . The primary master continues
19 to control the bus until time t_3 when it tristates the TxEnb line
20 of the UTOPIA bus. The clock cycle starting at t_4 is a "dead
21 cycle" during which neither bus master controls the bus and all
22 signals are held in the high state. The secondary master takes
23 control of the bus at time t_5 . Selection of the PHY device to
24 receive a cell from the secondary master is effected at time t_6 .

1 During cell transfer, the secondary master continues to poll the
2 address bus from the last address polled by the primary master.
3 Unless the secondary master re-asserts the Request line and
4 receives a Grant signal from the primary master before completing
5 cell transfer, control of the bus is returned to the primary
6 master when the secondary master completes cell transfer.

7
8 The presently preferred embodiment utilizes an arbitration
9 scheme to assure fairness in allocating control of the bus. The
10 presently preferred arbitration scheme requires that both bus
11 masters maintain a "scoreboard" of all PHYs indicating (with one
12 bit for each PHY) the results of the most recent poll (i.e. which
13 PHYs asserted the CLAV line). Each time a PHY is selected by
14 either master, the corresponding CLAV bit in the scoreboard is
15 reset. Neither master may select a PHY unless the corresponding
16 CLAV bit is set in the scoreboard. When both masters need to use
17 the bus, the primary master alternates control of the bus with the
18 secondary master.

19
20 Referring now to Figure 5, an exemplary apparatus 100
21 embodying the invention is based on the Transwitch® Cubit® family
22 of devices. The apparatus 100 has an "inlet side" and a "switch
23 side". The inlet side comprises a UTOPIA interface which
24 includes a cell inlet port 102 for receiving cell from a UTOPIA

1 bus and a cell outlet port 104 for transmitting cells to the
2 UTOPIA bus. As seen in Figure 5, the aforementioned Ready line 24
3 is illustrated as part (RxRDY) of the cell inlet port 102. The
4 aforementioned Request 26 and Grant 28 lines are shown as part
5 (TxReq and TxGnt) of the cell outlet port 104. These three
6 control lines are illustrated in Figure 5 with bidirectional
7 arrows because the apparatus 100 can be configured to act as
8 either a primary or secondary master. The outlet side of the
9 apparatus comprises a CellBus® interface 106 which includes
10 interface logic 108 and an arbiter and frame pulse generator 110.
11 A number of buffers, multiplexers and demultiplexers control the
12 flow of traffic between the inlet side and the switch side. In
13 addition, a number of interfaces to other external devices are
14 provided.

15
16 The Cell Inlet Port block 102 may be set to be compatible
17 with either UTOPIA Level 1 or Level 2. Cells from the UTOPIA bus
18 pass through a synchronization FIFO 112 before passing on to a
19 data queue 114. Optionally, cells may first pass through a
20 translation control 116 before entering the data queue 114.
21 Translation and routing header tables to support this function are
22 contained in an external static RAM (not shown). According to the
23 invention, the screening of cells may be performed in the
24 translation control 116 or in a block (not shown) between the FIFO

1 112 and the data queue 114. The data cells in the queue 114 are
2 multiplexed with control information in queues 118, 120 via
3 multiplexer 122 and passed to the CellBus® interface 108.
4

5 Cells received via the CellBus® interface 108 pass through a
6 cell address screen 124 which screens out cells having improper
7 addresses. Accepted cells are routed via demultiplexer 126 to
8 either a data queue 128, a loopback queue 130, or a control queue
9 132. The outlet data FIFO structure is configured in external
10 SSRAM (not shown) which interfaces with the device 100 via the
11 interface 134. Data cells from the queue 128 pass through the
12 SSRAM to an outlet synchronization queue 136 before passing to the
13 outlet port 104 onto the UTOPIA bus. Loopback cells from the
14 queue 130 are transferred to the queue 120 for transmission back
15 through the CellBus® interface 108. Control cells pass from the
16 queue 132 to an external microprocessor (not shown) via the
17 interface 138. Control cells returning from the microprocessor
18 are passed to the control queue 118 for transmission back through
19 the CellBus® interface 108.
20

21 The preferred embodiment of the device 100 also includes a
22 test access port 140 which conforms to the IEEE 1149.1 standard.

1 Referring now to Figure 6, an exemplary application of the
2 apparatus of the invention is shown in conjunction with the
3 Transwitch Phast™-12 SONET OC-12 uplink interface as well as other
4 network interfaces. As shown in Figure 6, two Phast™-3P SONET
5 OC-3 uplink interfaces 202, 204 are coupled to a UTOPIA bus 206.
6 The interface 202 is coupled to a first SONET OC-12 network 208
7 and the interface 204 is coupled to a second SONET OC-12 network
8 210. A first CUBIT® device 100a according to the invention is
9 coupled to the UTOPIA bus 206 and to a first bus 212 run according
10 to a CellBus® frame as described in U.S. Patents 6,104,724 and
11 5,901,146, the complete disclosures of which are hereby
12 incorporated by reference herein. A second CUBIT® device 100b
13 according to the invention is coupled to the UTOPIA bus 206 and to
14 a second bus 214. The first and second devices 100a, 110b of the
15 invention are coupled to each other as described above with
16 reference to Figure 3; one of them is designated the primary bus
17 master; and the other is designated the secondary bus master. A
18 third CUBIT® device 100c according to the invention is coupled to
19 the first bus 212 and to a frame relay network 216 via a high
20 level data link control 218. A fourth CUBIT® device 100d
21 according to the invention is coupled to the second bus 214 and to
22 an end user 220 via a digital subscriber line 222.

1 The exemplary arrangement shown in Figure 6 permits two OC-12
2 networks to be coupled via a single UTOPIA bus to two other
3 networks via two buses. The provision of two bus masters 100a,
4 100b on the UTOPIA bus 206 allows either OC-12 network 208, 210 to
5 communicate with either of the other networks 216, 220 and vice
6 versa.

7
8 There have been described and illustrated herein several
9 embodiments of methods and apparatus for supporting multiple
10 UTOPIA bus masters on a single UTOPIA bus. While particular
11 embodiments of the invention have been described, it is not
12 intended that the invention be limited thereto, as it is intended
13 that the invention be as broad in scope as the art will allow and
14 that the specification be read likewise. Thus, while examples of
15 the invention have been disclosed with reference to a UTOPIA-
16 CellBus® interface, it will be appreciated that invention could be
17 applied to any ATM layer device interface with a UTOPIA bus.
18 Also, while an exemplary implementation has been illustrated with
19 regard to OC-12 networks, etc., it will be appreciated that the
20 invention can be applied to any UTOPIA bus application where
21 multiple bus masters are desirable. Further, while the invention
22 has been disclosed with reference to three separate lines for
23 ready, request, and grant, fewer lines could be used for all three
24 signals.

TRA-065

1 It will therefore be appreciated by those skilled in the art
2 that yet other modifications could be made to the provided
3 invention without deviating from its spirit and scope as so
4 claimed.